

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: • The ACM Digital Library • The Guide

page <near/10> id* and determine <sentence> (modification

SEARCH

THE ACM DICITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used page near/10 id and determine sentence modification or change

Found **91,254** of **161,645**

Sort results by

Display

results

 Save results to a Binder

Search Tips

Open results in a new

window

Try an <u>Advanced Search</u>
Try this search in <u>The ACM Guide</u>

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10 nex

Best 200 shown

Relevance scale 🗆 🗖 🖬 🔳

1 Special issue: Al in engineering

D. Sriram, R. Joobbani

January 1985 ACM SIGART Bulletin, Issue 91

Full text available: pdf(8.79 MB)

Additional Information: full citation, abstract

The papers in this special issue were compiled from responses to the announcement in the July 1984 issue of the SIGART newsletter and notices posted over the ARPAnet. The interest being shown in this area is reflected in the sixty papers received from over six countries. About half the papers were received over the computer network.

2 Special issue on knowledge representation

Ronald J. Brachman, Brian C. Smith

February 1980 ACM SIGART Bulletin, Issue 70

Full text available: pdf(13.13 MB) Additional Information: full citation, abstract

In the fall of 1978 we decided to produce a special issue of the SIGART Newsletter devoted to a survey of current knowledge representation research. We felt that there were twe useful functions such an issue could serve. First, we hoped to elicit a clear picture of how people working in this subdiscipline understand knowledge representation research, to illuminate the issues on which current research is focused, and to catalogue what approaches and techniques are currently being developed. Secon ...

³ Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research

Full text available: pdf(4.21 MB)

Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

4 <u>Automated techniques for managing collections: Managing distributed collections: evaluating web page changes, movement, and replacement</u>



Zubin Dalal, Suvendu Dash, Pratik Dave, Luis Francisco-Revilla, Richard Furuta, Unmil Karadkar, Frank Shipman

June 2004 Proceedings of the 4th ACM/IEEE-CS joint conference on Digital libraries

Full text available: pdf(329.43 KB)

Additional Information: full citation, abstract, references, citings, index

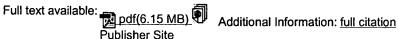
Distributed collections of Web materials are common. Bookmark lists, paths, and catalogs such as Yahoo! Directories require human maintenance to keep up to date with changes to the underlying documents. The Walden's Paths Path Manager is a tool to support the maintenance of distributed collections. Earlier efforts focused on recognizing the type and degree of change within Web pages and identifying pages no longer accessible. We now extend this work with algorithms for evaluating drastic changes ...

Keywords: change detection, collection management, document location

5 The FINITE STRING Newsletter: Abstracts of current literature

Computational Linguistics Staff

January 1987 Computational Linguistics, Volume 13 Issue 1-2



Specification and dialogue control of visual interaction through visual rewriting systems P. Bottoni, M. F. Costabile, P. Mussio



Full text available: pdf(886.71 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

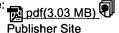
Computers are increasingly being seen not only as computing tools but more so as communication tools, thus placing special emphasis on human-computer interaction (HCI). In this article, the focus is on visual HCI, where the messages exchanged between human and computer are images appearing on the computer screen, as usual in current popular user interfaces. We formalize interactive sessions of a human-computer dialogue as a structured set of legal visual sentences, i.e., as a visual languag ...

Keywords: control automaton, dialogue control, visual languages

7 Speech repairs, intonational phrases, and discourse markers: modeling speakers' utterances in spoken dialogue

Peter A. Heeman, James F. Allen

December 1999 Computational Linguistics, Volume 25 Issue 4



Full text available: pdf(3.03 MB) Additional Information: full citation, abstract, references, citings

Interactive spoken dialogue provides many new challenges for natural language understanding systems. One of the most critical challenges is simply determining the speaker's intended utterances: both segmenting a speaker's turn into utterances and determining the intended words in each utterance. Even assuming perfect word recognition, the latter problem is complicated by the occurrence of speech repairs, which occur where speakers go back and change (or repeat) something they just said. The word ...

Special issue on natural language generation: Generating natural language summaries I





from multiple on-line sources

Dragomir R. Radev, Kathleen R. McKeown

September 1998 Computational Linguistics, Volume 24 Issue 3

Full text available: pdf(2.36 MB) Publisher Site

Additional Information: full citation, abstract, references, citings

We present a methodology for summarization of news about current events in the form of briefings that include appropriate background (historical) information. The system that we developed, SUMMONS, uses the output of systems developed for the DARPA Message Understanding Conferences to generate summaries of multiple documents on the same or related events, presenting similarities and differences, contradictions, and generalizations among sources of information. We describe the various components ...

9 Document Formatting Systems: Survey, Concepts, and Issues

Richard Furuta, Jeffrey Scofield, Alan Shaw

September 1982 ACM Computing Surveys (CSUR), Volume 14 Issue 3

Full text available: pdf(5.36 MB)

Additional Information: full citation, references, citings, index terms

10 Automatic parsing for content analysis

Frederick J. Damerau

June 1970 Communications of the ACM, Volume 13 Issue 6

Full text available: pdf(4.07 MB) Additional Information: full citation, abstract, references, citings

Although automatic syntactic and semantic analysis is not yet possible for all of an unrestricted natural language text, some applications, of which content analysis is one, do not have such a stringent coverage requirement. Preliminary studies show that the Harvard Syntactic Analyzer can produce correct and unambiguous identification of the subject and object of certain verbs for approximately half of the relevant occurences. This provides a degree of coverage for content analysis variable ...

Keywords: content analysis, information retrieval, language analysis, natural language processing, parsing, syntactic analysis, text processing

11 Special issue on natural language generation: Collaborative response generation in planning dialogues

Jennifer Chu-Carroll, Sandra Carberry

September 1998 Computational Linguistics, Volume 24 Issue 3

Full text available: pdf(3.45 MB) Additional Information: full citation, abstract, references, citings

In collaborative planning dialogues, the agents have different beliefs about the domain and about each other; thus, it is inevitable that conflicts arise during the planning process. In this paper, we present a plan-based model for response generation during collaborative planning, based on a recursive Propose-Evaluate-Modify framework for modeling collaboration. We focus on identifying strategies for content selection when 1) the system initiates information-sharing to gather furt ...

12 Abstract state machines capture parallel algorithms

Andreas Blass, Yuri Gurevich

October 2003 ACM Transactions on Computational Logic (TOCL), Volume 4 Issue 4

Full text available: pdf(610.28 KB) Additional Information: full citation, abstract, references, index terms

We give an axiomatic description of parallel, synchronous algorithms. Our main result is that every such algorithm can be simulated, step for step, by an abstract state machine with a background that provides for multisets.

Keywords: ASM thesis, Parallel algorithm, abstract state machine, postulates for parallel computation

13 Spoken dialogue technology: enabling the conversational user interface

Michael F. McTear

March 2002 ACM Computing Surveys (CSUR), Volume 34 Issue 1

Full text available: pdf(987.69 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

Spoken dialogue systems allow users to interact with computer-based applications such as databases and expert systems by using natural spoken language. The origins of spoken dialogue systems can be traced back to Artificial Intelligence research in the 1950s concerned with developing conversational interfaces. However, it is only within the last decade or so, with major advances in speech technology, that large-scale working systems have been developed and, in some cases, introduced into commerc ...

Keywords: Dialogue management, human computer interaction, language generation, language understanding, speech recognition, speech synthesis

14 Unsupervised learning of the morphology of a natural language

John Goldsmith

June 2001 Computational Linguistics, Volume 27 Issue 2

Publisher Site

Full text available: pdf(3.19 MB) Additional Information: full citation, abstract, references, citings

This study reports the results of using minimum description length (MDL) analysis to model unsupervised learning of the morphological segmentation of European languages, using corpora ranging in size from 5,000 words to 500,000 words. We develop a set of heuristics that rapidly develop a probabilistic morphological grammar, and use MDL as our primary tool to determine whether the modifications proposed by the heuristics will be adopted or not. The resulting grammar matches well the analysis that ...

¹⁵ Parallel execution of prolog programs: a survey

Gopal Gupta, Enrico Pontelli, Khayri A.M. Ali, Mats Carlsson, Manuel V. Hermenegildo July 2001 ACM Transactions on Programming Languages and Systems (TOPLAS), Volume 23 Issue 4

Full text available: pdf(1.95 MB)

Additional Information: full citation, abstract, references, citings, index terms

Since the early days of logic programming, researchers in the field realized the potential for exploitation of parallelism present in the execution of logic programs. Their high-level nature, the presence of nondeterminism, and their referential transparency, among other characteristics, make logic programs interesting candidates for obtaining speedups through parallel execution. At the same time, the fact that the typical applications of logic programming frequently involve irregular computatio ...

Keywords: Automatic parallelization, constraint programming, logic programming, parallelism, prolog

16

An empirically based system for processing definite descriptions

Renata Vieira, Massimo Poesio

December 2000 Computational Linguistics, Volume 26 Issue 4

Publisher Site

Full text available: pdf(2.94 MB) Additional Information: full citation, abstract, references

We present an implemented system for processing definite descriptions in arbitrary domains. The design of the system is based on the results of a corpus analysis previously reported, which highlighted the prevalence of discourse-new descriptions in newspaper corpora. The annotated corpus was used to extensively evaluate the proposed techniques for matching definite descriptions with their antecedents, discourse segmentation, recognizing discourse-new descriptions, and suggesting anchors for brid ...

17 Building efficient and effective metasearch engines

Weiyi Meng, Clement Yu, King-Lup Liu

March 2002 ACM Computing Surveys (CSUR), Volume 34 Issue 1

Full text available: pdf(416.07 KB)

Additional Information: full citation, abstract, references, citings, index terms

Frequently a user's information needs are stored in the databases of multiple search engines. It is inconvenient and inefficient for an ordinary user to invoke multiple search engines and identify useful documents from the returned results. To support unified access to multiple search engines, a metasearch engine can be constructed. When a metasearch engine receives a query from a user, it invokes the underlying search engines to retrieve useful information for the user. Metasearch engines have ...

Keywords: Collection fusion, distributed collection, distributed information retrieval, information resource discovery, metasearch

18 Generating summaries of multiple news articles

Kathleen McKeown, Dragomir R. Radev

July 1995 Proceedings of the 18th annual international ACM SIGIR conference on Research and development in information retrieval

Full text available: pdf(1.02 MB)

Additional Information: full citation, references, citings, index terms

Keywords: natural language generation, natural language summarization, summarization of multiple texts

19 A stochastic finite-state word-segmentation algorithm for Chinese

Richard Sproat, William Gale, Chilin Shih, Nancy Chang

September 1996 Computational Linguistics, Volume 22 Issue 3

Full text available:

Publisher Site

pdf(1.91 MB) Additional Information: full citation, abstract, references, citings

The initial stage of text analysis for any NLP task usually involves the tokenization of the input into words. For languages like English one can assume, to a first approximation, that word boundaries are given by whitespace or punctuation. In various Asian languages, including Chinese, on the other hand, whitespace is never used to delimit words, so one must resort to lexical information to "reconstruct" the word-boundary information. In this paper we present a stochastic finite-state model whe ...

20

The FINITE STRING newsletter: Abstracts of current literature

American Journal of Computational Linguistics Staff July 1981 Computational Linguistics, Volume 7 Issue 3

Full text available: pdf(2.42 MB) Publisher Site

Additional Information: full citation

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright @ 2005 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

+"database*" and +"Isn" < near/10 > page < near/10 > id* and

SEARCH



Feedback Report a problem Satisfaction survey

Terms used database and Isn near/10 page near/10 id and determine sentence modification

Found 33 of 161,645

Sort results by Display

Results 1 - 20 of 33

relevance

Save results to a Binder Search Tips

Try an Advanced Search Try this search in The ACM Guide

expanded form Open results in a new results window

Result page: 1 2

Relevance scale

1 ARIES: a transaction recovery method supporting fine-granularity locking and partial rollbacks using write-ahead logging

C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz March 1992 ACM Transactions on Database Systems (TODS), Volume 17 Issue 1

Full text available: pdf(5.23 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

DB2TM, IMS, and TandemTM systems. ARIES is applicable not only to database management systems but also to persistent object-oriented languages, recoverable file systems and transaction-based operating systems. ARIES has been implemented, to varying degrees, in IBM's OS/2TM Extended Edition Database Manager, DB2, Workstation Data Save Facility/VM, Starburst and QuickSilver, and in the University of Wisconsin's EXODUS and Gamma d ...

Keywords: buffer management, latching, locking, space management, write-ahead logging

2 Concurrency control and recovery for balanced B-link trees

Ibrahim Jaluta, Seppo Sippu, Eljas Soisalon-Soininen

April 2005 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 14 Issue 2

Full text available: pdf(302.02 KB) Additional Information: full citation, abstract

In this paper we present new concurrent and recoverable B-link-tree algorithms. Unlike previous algorithms, ours maintain the balance of the B-link tree at all times, so that a logarithmic time bound for a search or an update operation is guaranteed under arbitrary sequences of record insertions and deletions. A database transaction can contain any number of operations of the form "fetch the first (or next) matching record", "insert a record", or "delete a reco ...

Keywords: Concurrency control, Recovery, Transaction, Tree-structure modifications

3 Virtual memory management for database systems Irving L. Traiger



October 1982 ACM SIGOPS Operating Systems Review, Volume 16 Issue 4

Full text available: pdf(2.08 MB) Additional Information: full citation, abstract, references, citings

Over the last several years, a number of hardware and software systems have been developed which map entire files directly into the virtual memory address spaces used by programs. Since all file contents are directly addressable, there is no need for a programmer to issue explicit file system actions, such as Read or Write. In addition, all of the buffer management problems are eliminated, since programmers do not have to squeeze pieces of large files into small virtual spaces. Although these ad ...

4 Crash recovery in client-server EXODUS

Michael J. Franklin, Michael J. Zwilling, C. K. Tan, Michael J. Carey, David J. DeWitt June 1992 ACM SIGMOD Record, Proceedings of the 1992 ACM SIGMOD international conference on Management of data, Volume 21 Issue 2

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(1.50 MB)

In this paper, we address the correctness and performance issues that arise when implementing logging and crash recovery in a page-server environment. The issues result from two characteristics of page-server systems: 1) the fact that data is modified and cached in client database buffers that are not accessible by the server, and 2) the performance and cost trade-offs that are inherent in a client-server environment. We describe a recovery system that we have implemented for the client-ser ...

⁵ Model and verification of a data manager based on ARIES

Dean Kuo

December 1996 ACM Transactions on Database Systems (TODS), Volume 21 Issue 4

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(813.93 KB) terms, review

In this article, we model and verify a data manager whose algorithm is based on ARIES. The work uses the I/O automata method as the formal model and the definition of correctness is defined on the interface between the scheduler and the data manager.

Keywords: ARIES, I/O automata, system failures

6 Concurrency and recovery in generalized search trees

Marcel Kornacker, C. Mohan, Joseph M. Hellerstein

June 1997 ACM SIGMOD Record, Proceedings of the 1997 ACM SIGMOD international conference on Management of data, Volume 26 Issue 2

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(1.59 MB)

This paper presents general algorithms for concurrency control in tree-based access methods as well as a recovery protocol and a mechanism for ensuring repeatable read. The algorithms are developed in the context of the Generalized Search Tree (GiST) data structure, an index structure supporting an extensible set of queries and data types. Although developed in a GiST context, the algorithms are generally applicable to many treebased access methods. The concurrency control protocol is base ...

⁷ Efficient transaction support for dynamic information retrieval systems Mohan Kamath, Krithi Ramamritham

August 1996 Proceedings of the 19th annual international ACM SIGIR conference on Research and development in information retrieval

Full text available: 📆 pdf(1.09 MB) Additional Information: full citation, references, citings, index terms





















8 Industrial sessions: beyond relational tables: Coordinating backup/recovery and data consistency between database and file systems



Suparna Bhattacharya, C. Mohan, Karen W. Brannon, Inderpal Narang, Hui-I Hsiao, Mahadevan Subramanian

June 2002 Proceedings of the 2002 ACM SIGMOD international conference on Management of data

Full text available: 📆 pdf(1.44 MB)

Additional Information: full citation, abstract, references, index terms

Managing a combined store consisting of database data and file data in a robust and consistent manner is a challenge for database systems and content management systems. In such a hybrid system, images, videos, engineering drawings, etc. are stored as files on a file server while meta-data referencing/indexing such files is created and stored in a relational database to take advantage of efficient search. In this paper we describe solutions for two potentially problematic aspects of such a data ...

Keywords: DB2, content management, database backup, database recovery, datalinks

ARIES/CSA: a method for database recovery in client-server architectures
 C. Mohan, Inderpal Narang



May 1994 ACM SIGMOD Record , Proceedings of the 1994 ACM SIGMOD international conference on Management of data, Volume 23 Issue 2

Full text available: pdf(1.33 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

This paper presents an algorithm, called ARIES/CSA (Algorithm for Recovery and Isolation Exploiting Semantics for Client-Server Architectures), for performing recovery correctly in client-server (CS) architectures. In CS, the server manages the disk version of the database. The clients, after obtaining database pages from the server, cache them in their buffer pools. Clients perform their updates on the cached pages and produce log records. The log records are buffered loca ...

10 ARIES/IM: an efficient and high concurrency index management method using writeahead logging



C. Mohan, Frank Levine

June 1992 ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data, Volume 21 Issue 2

Full text available: pdf(1.32 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> <u>terms</u>

This paper provides a comprehensive treatment of index management in transaction systems. We present a method, called ARIESIIM (Algorithm for Recovery and Isolation Exploiting Semantics for Index Management), for concurrency control and recovery of B+trees. ARIES/IM guarantees serializability and uses write-ahead logging for recovery. It supports very high concurrency and good performance by (1) treating as the lock of a key the same lock as the one on the ...

11 Atomic incremental garbage collection and recovery for a large stable heap Elliot K. Kolodner, William E. Weihl



June 1993 ACM SIGMOD Record, Proceedings of the 1993 ACM SIGMOD international conference on Management of data, Volume 22 Issue 2

Full text available: pdf(1.34 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

A stable heap is storage that is managed automatically using garbage collection,

manipulated using atomic transactions, and accessed using a uniform storage model. These features enhance reliability and simplify programming by preventing errors due to explicit deallocation, by masking failures and concurrency using transactions, and by eliminating the distinction between accessing temporary storage and permanent storage. Stable heap management is useful for programming lang ...

12 Garbage collection for a client-server persistent object store

Laurent Amsaleg, Michael J. Franklin, Olivier Gruber

August 1999 ACM Transactions on Computer Systems (TOCS), Volume 17 Issue 3

Full text available: pdf(267.18 KB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms, review

We describe an efficient server-based algorithm for garbage collecting persistent object stores in a client-server environment. The algorithm is incremental and runs concurrently with client transactions. Unlike previous algorithms, it does not hold any transactional locks on data and does non require callbacks to clients. It is fault-tolerant, but performs very little logging. The algorithm has been designed to be integrated into existing systems, and therefore it works with standard i ...

Keywords: client-server system, logging, persistent object-store, recovery

13 <u>Multi-level transaction management for complex objects: implementation, performance, parallelism</u>

Gerhard Weikum, Christof Hasse

October 1993 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 2 Issue 4

Full text available: pdf(2.83 MB) Additional Information: full citation, abstract, references, citings

Multi-level transactions are a variant of open-nested transactions in which the subtransactions correspond to operations at different levels of a layered system architecture. They allow the exploitation of semantics of high-level operations to increase concurrency. As a consequence, undoing a transaction requires compensation of completed subtransactions. In addition, multi-level recovery methods must take into consideration that high-level operations are not necessarily atomic if multiple pages ...

Keywords: atomicity, complex objects, inter- and intratransaction parallelism, multi-level transactions, performance, persistence, recovery

14 Logical logging to extend recovery to new domains

David Lomet, Mark Tuttle

June 1999 ACM SIGMOD Record, Proceedings of the 1999 ACM SIGMOD international conference on Management of data, Volume 28 Issue 2

Full text available: pdf(1.65 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

Recovery can be extended to new domains at reduced logging cost by exploiting "logical" log operations. During recovery, a logical log operation may read data values from any recoverable object, not solely from values on the log or from the updated object. Hence, we needn't log these values, a substantial saving. In [8], we developed a redo recovery theory that deals with general log operations and proved that the stable database remains recoverable when it is explained in terms ...

15 Analysis of recovery in a database system using a write-ahead log protocol Anant Jhingran, Pratap Khedkar

June 1992 ACM SIGMOD Record, Proceedings of the 1992 ACM SIGMOD international

conference on Management of data, Volume 21 Issue	conference	e on Managem	ent of data.	Volume 21	Issue 2
---	------------	--------------	--------------	-----------	---------

Full text available: pdf(875.06 KB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

In this paper we examine the recovery time in a database system using a Write-Ahead Log protocol, such as ARIES [9], under the assumption that the buffer replacement policy is strict LRU. In particular, analytical equations for log read time, data I/O, log application, and undo processing time are presented. Our initial model assumes a read/write ratio of one, and a uniform access pattern. This is later generalized to include different read/write ratios, as well as a "hot set" m ...

16 Parallelism in relational data base systems: architectural issues and design approaches

Hamid Pirahesh, C. Mohan, Josephine Cheng, T. S. Liu, Pat Selinger

July 1990 Proceedings of the second international symposium on Databases in parallel and distributed systems

Full text available: 1 pdf(2.50 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

With current systems, some important complex queries may take days to complete because of: (1) the volume of data to be processed, (2) limited aggregate resources. Introducing parallelism addresses the first problem. Cheaper, but powerful computing resources solve the second problem. According to a survey by Brodie,1 only 10% of computerized data is in data bases. This is an argument for both more variety and volume of data to be moved into data base systems. We conject ...

17 Recovery guarantees for Internet applications

Roger Barga, David Lomet, German Shegalov, Gerhard Weikum August 2004 **ACM Transactions on Internet Technology (TOIT)**, Volume 4 Issue 3

Full text available: pdf(997.52 KB) Additional Information: full citation, abstract, references, index terms

Internet-based e-services require application developers to deal explicitly with failures of the underlying software components, for example web servers, servlets, browser sessions, and so forth. This complicates application programming, and may expose failures to end users. This paper presents a framework for an application-independent infrastructure that provides recovery guarantees and masks almost all system failures, thus relieving the application programmer from having to deal with these f ...

Keywords: Exactly-once execution, application recovery, communication protocols, interaction contracts

18 Efficient and flexible methods for transient versioning of records to avoid locking by read-only transactions

C. Mohan, Hamid Pirahesh, Raymond Lorie

June 1992 ACM SIGMOD Record, Proceedings of the 1992 ACM SIGMOD international conference on Management of data, Volume 21 Issue 2

Full text available: pdf(1.19 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> <u>terms</u>

We present efficient and flexible methods which permit read-only transactions that do not mind reading a possibly slightly old, but still consistent, version of the data base to execute without acquiring locks. This approach avoids the undesirable interferences between such queries and the typically shorter update transactions that cause unnecessary and costly delays. Indexed access by such queries is also supported, unlike by the earlier methods. Old versions of records are maintained only ...

19 Research papers: storage, indexing, and system architecture: Online B-tree merging Xiaowei Sun, Rui Wang, Betty Salzberg, Chendong Zou



June 2005 Proceedings of the 2005 ACM SIGMOD international conference on Management of data

Full text available: pdf(394.41 KB) Additional Information: full citation, abstract, references

Many scenarios involve merging of two B-tree indexes, both covering the same key range. Increasing demand for continuous availability and high performance requires that such merging be done online, with minimal interference to normal user transactions. In this paper we present an online B-tree merging method, in which the merging of leaf pages in two B-trees are piggybacked lazily with normal user transactions, thus making the merging I/O efficient and allowing user transactions to access only o ...

20 On-line reorganization of sparsely-populated B+-trees

Chendong Zou, Betty Salzberg

June 1996 ACM SIGMOD Record, Proceedings of the 1996 ACM SIGMOD international conference on Management of data, Volume 25 Issue 2

Full text available: pdf(1.17 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

In this paper, we present an efficient method to do online reorganization of sparsely-populated B⁺-trees. It reorganizes the leaves first, compacting in short operations groups of leaves with the same parent. After compacting, optionally, the new leaves may swap locations or be moved into empty pages so that they are in key order on the disk. After the leaves are reorganized, the method shrinks the tree by making a copy of the upper part of the tree while leaving the leaves in place. ...

Results 1 - 20 of 33 Result page: 1 2 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

<u>Terms of Usage Privacy Policy Code of Ethics Contact Us</u>

Useful downloads: Adobe Acrobat Q QuickTime Windows Media Player Real Player